

# Set-merging for the Matching Algorithm of Micali and Vazirani

Harold N. Gabow\*

July 4, 2013

## Abstract

The algorithm of Micali and Vazirani [4] finds a maximum cardinality matching in time  $O(\sqrt{nm})$  if an efficient set-merging algorithm is used. The latter is provided by the incremental-tree set-merging algorithm of [2]. Details of this application to matching were omitted from [2] and are presented in this note.

## 1 Introduction

Micali and Vazirani implemented the “shortest augmenting path” approach to find a maximum cardinality matching in time  $O(\sqrt{nm})$  [4]. This time bound assumes a linear-time algorithm for set-merging is used. It is tempting to conjecture (as was done in [4]) that the special structure of blossoms achieves this goal automatically. However this claim remains unsubstantiated. Instead the incremental-tree set-merging algorithm of [2] can be used to represent blossoms and achieve the desired time bound. The details of this application to matching were omitted from [2] and they are presented in this note. For accuracy we use the recent paper [5] as the source of all details for the cardinality matching algorithm. We assume familiarity with that paper.<sup>1</sup>

We will present an implementation of the Micali-Vazirani algorithm that achieves time  $O(m)$  per phase on a RAM. This easily implies the desired time bound  $O(\sqrt{nm})$ . The issue is the use of set-merging to represent blossoms. Towards this end we divide each phase into two subphases. The first subphase assigns levels to vertices. In this subphase all free vertices are essentially equivalent. This allows the search structure to be modeled by a tree, and allows use of incremental-tree set merging to model blossoms. The second subphase finds augmenting paths. Here we choose bridges for DDFS in a depth-first manner. This gives a linear order to newly created blossoms, allowing set merging to be handled with a stack. (This stack-based set-merging is used in [1] to compute strong components efficiently.) The first subphase requires a RAM, the second subphase runs on both RAM and pointer machine.

The rest of this section gives background from [5]. Then the set-merging algorithm is presented, the first “assignment” subphase in Section 2 and the second “augmentation” subphase in Section 3.

The entire discussion uses this terminology: A *vertex* is a vertex of the given graph  $G$ . A *blossom* is a blossom from the previous search level (i.e., tenacity  $< 2i + 1$  during search level  $i$ ); a *base* is the base vertex of such a blossom, or a vertex of finite evenlevel and empty blossom. Thus

---

\*Department of Computer Science, University of Colorado at Boulder, Boulder, CO 80309. E-mail: [hal@cs.colorado.edu](mailto:hal@cs.colorado.edu)

<sup>1</sup>Another cardinality matching algorithm running in time  $O(\sqrt{nm})$  is presented in [3]. The use of incremental-tree set-merging in that algorithm is straightforward.

at the start of a search level, the vertices of  $G$  with finite evenlevel are partitioned into blossoms and bases.

As in [5, Definition 29] every vertex  $v$  has a value  $\text{bud}^*(v)$ . A *petal\** is the set of all vertices with finite evenlevel that have the same  $\text{bud}^*$ . Assume a blossom has  $\text{bud}$  equal to its base, so when a search level begins, each blossom plus its base is a *petal\**. At any point in time the *petal\**'s partition the set of vertices with finite evenlevel. A *free petal\** is a *petal\** whose  $\text{bd}^*$  is a free vertex. We say  $\text{bud}^*(P) = b$  if  $P$  is a *petal\** corresponding to  $\text{bud}^* b$ .

Note that  $\text{bud}^*(v) = v$  iff  $v$  is the bud of a *petal\**, or (as in [5])  $v$  is not in any *petal\**.

The following lemma is adapted from [5]. We will use parts (i) and (iii); we include (ii) for completeness.

**Lemma 1.1** *Consider an arbitrary  $\text{petal}^* P^*$  with bud  $b$ . Let  $v$  be an arbitrary vertex in  $P^*$ .*

(i) *Every  $\text{minlevel}(v)$  path  $p$  goes through  $b$ . Its subpath from  $v$  to  $b$ , denoted  $\text{minlevel}(v; b)$  (like the notation of [5]) is contained in  $P^*$ .*

(ii) *The same holds for a  $\text{maxlevel}(v)$  path  $p$  whose (unique)  $\text{tenacity}(v)$  bridge has been processed by MAX.*

(iii)  *$P^*$  is precisely equal to the union of the paths  $\text{minlevel}(v, b)$ ,  $v \in P^*$ . In fact we can restrict the union to vertices  $v$  in the support of  $P^*$ , assuming we add in the *petal\** whose  $\text{bud}^*$  is  $v$ .*

**Remarks:** Note that (ii) need not hold for a  $\text{maxlevel}$  path using a bridge that has not been processed by MAX (i.e., the *petal* has not been expanded to its complete blossom).

Also note that parts (i)–(ii) imply the Micali-Vazirani algorithm handles augmenting paths correctly: Suppose  $P^*$  of the lemma is on an augmenting path; clearly it contains  $b$ . Suppose for contradiction that  $v \in P^*$  is on another ap (ap = shortest augmenting path)  $Q$  that does not contain  $b$ . Part (i) shows  $Q$  cannot contain a  $\text{minlevel}(v)$  path. [5, Theorem 26] shows a  $\text{maxlevel}(v)$  path  $M$  in  $Q$  uses a bridge of  $\text{tenacity} < l_m$ . So part (ii) applies to  $M$ . Thus  $M$  contains  $b$  and cannot be in  $Q$ .

**Proof:** Let  $t$  be the  $\text{tenacity}$  in the current search level.

Define a subgraph  $H(P^*)$ : Start with the graph in which  $P^*$  is formed (i.e., previous augmenting paths, if any, have been deleted). Then remove every bridge that has not been processed by MAX. Apply [5, Lemma 30] to  $P^*$ . If  $P^*$  has no  $\text{tenacity } t$  vertices the lemma shows  $P^* - b$  is a blossom of  $\text{tenacity } t - 2$ . In the opposite case  $P^* - b$  is a blossom of  $H(P^*)$  of  $\text{tenacity } t$ . Since  $H(P^*)$  has  $l_m > t$ , [5, Theorem 7] applies. It shows that as in parts (i) and (ii), each  $\text{minlevel}(v)$  and  $\text{maxlevel}(v)$  path  $M$  passes through  $b$ .

[5, Definition 17] of blossom, plus a simple induction, show  $M$  is contained in  $P^*$ . This completes the proof of (i) and (ii). It also proves (iii).  $\square$

## 2 Assignment Subphase

Each phase begins in the Assignment Subphase. This subphase assigns levels, according to the procedures MIN and MAX of [5, Algorithm 1]. The subphase ends when a DDFS discovers the first augmenting path. At that point the Augmentation Subphase begins.

The Assignment Subphase is based on a search forest  $F$ . Section 2.1 defines  $F$  and discusses its properties. Then Section 2.2 gives the algorithm for the subphase.

## 2.1 The Search Forest

At any point in a phase consider the following forest  $F$  (which represents the current search structure).  $F$  has a node for every  $\text{bd}^* v$  that has a finite minlevel. Each free petal\* is a root. (At the start of the phase this means each free vertex is a root.) The parent of  $v$  corresponds to the first predecessor of  $v$ . More precisely suppose MIN scans an edge  $(u, v)$  and changes  $\text{minlevel}(v)$  from  $\infty$  to a finite value. Here  $u$  is a vertex, and  $\text{bud}^*(u)$  is a node of  $F$ . The new tree  $F$  has  $\text{bud}^*(u)$  the parent of  $v$ . Thus the  $F$ -path from a node  $v$  (a  $\text{bud}^*$ ) to its root, say  $F(v)$ , corresponds (by contraction) to a  $\text{minlevel}(v)$  path of props from  $v$  to a free petal\*.

This completes the definition of  $F$ . Now we describe how  $F$  gets modified by the MIN and MAX procedures.

If MIN assigns a first predecessor to  $v$ ,  $F$  gets a new leaf  $v$ .

Suppose MAX forms a new petal  $P$  for bridge  $rg$  of tenacity  $t = 2i + 1$ . Let  $P^*$  be the corresponding petal\*, and let  $b$  be its bud (i.e., the bottleneck of the DDFS for  $rg$ ). Take any  $v \in P^*$  that is a  $\text{bd}^*$  at the start of the current DDFS. Lemma 1.1(i) shows the  $vb$ -subpath of  $F(v)$  is contained in  $P^*$ . So Lemma 1.1(iii) shows the entire petal  $P^*$  corresponds to a subtree  $S$  of  $F$ . We conclude the new  $F$  is formed by contracting the edges of  $S$  in  $F$ .

## 2.2 The Algorithm

We use incremental-tree set merging as the data structure for  $F$ . Let  $D$  be the forest of this data structure. A  $\text{bd}^* b$  in  $F$  corresponds to a subtree of  $D$  that has been merged into its root node  $b$ . Beyond that the two forests are identical.

The MIN step performs a *grow* operation on  $D$  each time a first predecessor is processed.

A MAX step will execute DDFS and either find the next petal\* or find the first augmenting path.

First suppose DDFS, say for bridge  $rg$ , identifies a petal\*  $P^*$  with bud  $b$ . The DDFS has identified the support of  $rg$  (in the trees  $R$  and  $G$  of [5]). The algorithm performs a *union* operation on  $D$  for each support node. This merges the entire petal  $P^*$  into its bud  $b$ . (To verify this we remark that a petal\*  $Q^*$  whose bud\*  $q$  is in the support of  $rg$  is contained in  $P^*$ , since every vertex of  $Q^*$  gets a new  $\text{bd}^* b$ . So the *unions* correctly contract  $P^*$ .) The updated  $D$  corresponds to the new search forest  $F$  by the description of MAX in Section 2.1.

It remains to describe how the DDFS is implemented, and in particular how it detects the first augmenting path. Assume the DDFS begins with the sets of  $D$  correctly corresponding to the current petal\*s. The DDFS uses *find* operations on  $D$  to jump from a vertex to its  $\text{bd}^* x$ ; it moves from  $x$  using a prop. Observe that the above procedure to update  $F$  and the data structure  $D$  occurs after the DDFS. So all the *find* operations in the DDFS give the correct bud\*. Thus the DDFS works correctly, discovering the next petal\* (which might be a free petal\*) or the first augmenting path.

## 2.3 Time Bound

The time for the Assignment Subphase is  $O(m)$ . In proof, the time for a Micali-Vazirani phase is  $O(m)$  plus the time for set merging. Incremental-tree set merging uses  $O(n + m)$  time for  $m$  *find* operations,  $n$  *make\_root* or *grow* operations, and  $n$  *union* operations.

### 3 Augmentation Subphase

The last search level  $i_m$  of the phase is implemented by the Augmentation Subphase. During this subphase the algorithm will know  $\text{bud}^*(v)$  for every vertex  $v$  in the following sense. At the start of the phase  $\text{bud}^*(v) = \text{base}^*(v)$  from phase  $i_{m-1}$ . At some point  $v$  may enter a search stack. It remains in the stack until it is deleted, either because of an ap or because it is known not to be in any ap. From this point on  $\text{bud}^*(v)$  will always be computed using the data structure for the stack.

The stack models this variant of DDFS: A *DDFS anchored on  $A$* , for bridge  $rg$ , starts with a (shortest) path  $A$  from level  $l_0$  to a vertex  $r$  which is on a bridge  $rg$  of tenacity  $l_m$ . Denote  $A$  as the sequence  $A_1, A_2, \dots, A_\ell$ , where each  $A_i$  is a vertex or a petal\* (a vertex  $A_i$  will be a vertex of finite oddlevel not in any petal\*). As usual the DDFS visits some set of vertices, say  $X$ , and finds either an ap or a new petal\* which actually is  $X$ . The anchored DDFS has this defining property: If an ap is found then  $V(A) \subseteq X$ . If  $X$  is a petal\*, say with bud  $b$ , then  $b = \text{bud}^*(A_i)$  for some  $i$ , and  $V(A) \cap X$  consists of the sets  $A_j, j \geq i$ .

We note that the algorithm will be valid for a pointer machine.

#### 3.1 Data Structure

Array  $\hat{A}$  stores the sequence  $V(A_1), \dots, V(A_\ell)$ . Each vertex  $v$  has a pointer (index)

$$a[v] = \begin{cases} 0 & v \notin V(A) \\ i & \hat{A}[i] = v. \end{cases}$$

A companion array gives the buds for  $A$ :

$$\hat{B}[i] = \text{bud}^*(A_i), \quad i = 1, \dots, \ell.$$

$\hat{A}$  stores  $\text{bud}^*(A_i)$  before the other vertices of  $A_i$ . Thus for a vertex of  $G$  in  $V(A)$ ,  $\text{bud}^*(v)$  is the vertex  $b = \hat{B}[i]$  where  $i$  is the maximum index satisfying  $a[b] \leq a[v]$ .

In addition there are some bookkeeping lists, that are easily maintained: Each petal\*  $P$  has a list of props  $uv$ ,  $u \in P$  and  $u$  a predecessor of  $v$ , and bridges  $uv$  of tenacity  $\ell_m$ ,  $u \in P$  ( $v$  may or may not be in  $P$ ). Each blossom (from  $i_{m-1}$ ) has a list of its  $G$ -vertices (it will be used to populate  $\hat{A}$ ).

#### 3.2 Algorithm

##### Highlevel Outline

$\hat{A}$  is managed as a stack. The main loop computes  $A$  and then does a DDFS anchored on  $A$ . If an ap is discovered all of  $V(A)$  gets deleted from the graph (by definition of anchored DDFS). If a bottleneck  $b$  is discovered,  $b = \text{bud}^*(A_i)$ , we can merge petal\*s  $A_i, \dots, A_\ell$  (again definition of anchored DDFS). So pop the  $\hat{B}$  stack to erase the buds above  $b$ , and add the new vertices of the petal\* to  $\hat{A}$ . Also merge appropriate bookkeeping lists for the new petal\*. Then do a dfs from  $A_i$ , using props, trying to complete a new  $A$ -path (i.e., search for a base or blossom that has an incident bridge). If one is found start the next iteration. (Do the same thing in the ap case.) If not,  $A_i$  cannot be on any ap. Delete  $A_i$  from the graph, and dfs search from  $A_{i-1}$ .

##### Detailed Algorithm

### 3.2.1 Computing $A$

In the following, each time  $A$  is extended by a new vertex or petal\* (which must be a blossom plus its base)  $A_i$ , push  $V(A_i)$  onto  $\hat{A}$  and push  $\text{bud}^*(A_i)$  onto  $\hat{B}$ .

We start with  $A$  as the sequence  $A_1, \dots, A_i$  ( $i = \epsilon$ ) from the previous search; if  $A$  is empty, initialize it to a free petal\*  $A_1$ , and set  $i = 1$ .

Do a dfs from  $A_i$ , using props to find the next vertex or petal\*. Extend the dfs path  $A$  (0 or more times) until it reaches a vertex or petal\*  $A_\ell$  that either is either a dead end or it has an incident bridge.

If  $A_\ell$  has no outgoing props, it cannot be in any ap. Delete  $A_\ell$  from  $A$  and the graph. Continue the dfs.

If  $A_\ell$  has an incident bridge  $rg$ , do a DDFS anchored on  $A$ .

### 3.2.2 Anchored DDFS

At any time in the DDFS say each DFS path ends at its *active vertex* (called “center of activity” in [5]). We will maintain the invariant that some vertex  $\alpha \in V(A)$  is active; more precisely there is always an index  $\epsilon$  such that  $\alpha = \text{bud}^*(A_\epsilon)$  is active and the DDFS has visited  $\text{bud}^*(A_i)$  iff  $i \geq \epsilon$ . Whenever the algorithm moves  $\alpha$ , it descends to the preceding vertex of  $A$ . Thus the invariant is preserved. The invariant implies that the DDFS correctly implements anchored DDFS (i.e., if an ap is found then  $V(A) \subseteq X$ ; if not then  $X$  is a petal\* with bud  $\text{bud}^*(A_\epsilon)$  for some  $\epsilon$ , and  $V(A) \cap X = \bigcup_{i \geq \epsilon} V(A_i)$ ). Let  $u$  denote the other active vertex (perhaps  $u = \alpha$ ).

The algorithm always records  $\epsilon$ . It moves  $\alpha$  by decrementing  $\epsilon$  by 1.

The rule for how the DDFS moves is determined by its state. The 2 possible states are called *Disjoint* (when  $u \notin V(A)$ ) and *Meet* (when  $u = \alpha$ ). (There may be  $> 1$  move between states, e.g., the algorithm starts backtracking when it enters state *Meet*.)

*State Disjoint*,  $u \notin V(A)$ .

If  $\text{level}(\alpha) \geq \text{level}(u)$ , move  $\alpha$  to the preceding vertex of  $A$ . Stay in state Disjoint and make the next move.

If  $\text{level}(u) > \text{level}(\alpha)$ , let  $uv$  be a prop from  $u$  to a lower level. (It exists since we’re not backtracking.) If  $a[v] = 0$ , move  $u$  to  $\text{bud}^*(v)$ . Stay in state Disjoint and make the next move.

If not (i.e.,  $a[v] \neq 0$ ) then  $v \in V(A)$ . If  $\epsilon > 1$  let  $b = \hat{B}[\epsilon - 1]$ , i.e., the bud preceding  $\alpha$  in  $A$ . If  $a[v] \geq a[b]$ ,  $\text{bud}^*(v)$  has been visited, so reject prop  $uv$ . Otherwise  $a[v] < a[b]$ , so  $a[v] \leq a[\alpha]$ . Move  $\alpha$  0 or more times until  $a[\alpha] \leq a[v]$ . Move  $u$  to  $\alpha$ . Now the 2 DFS paths have met so enter state Meet.

*State Meet*,  $u = \alpha$ .

Enter backtracking mode. Proceed as usual. Note that  $\alpha$  always remains active during the backtracking. As before the backtracking DFS always backs off from a previously visited vertex. It scans the next prop following the rule for state Disjoint, until (as usual) it finds a vertex  $u \neq \alpha$  at level  $\leq \text{level}(\alpha)$ . If  $u \notin V(A)$  it goes to state Disjoint. If  $u \in V(A)$ , it will reenter state Meet at the new, lower  $\alpha$ .

*Initializing the DDFS for bridge  $rg$ .*

Set  $\epsilon = \ell$ ,  $\alpha = \text{bud}^*(A_\epsilon)$ . If  $a[g] = 0$  then set  $u = \text{bud}^*(g)$  and begin the DDFS in state Disjoint. Otherwise  $g \in V(A)$ . If  $a[g] \geq a[\alpha]$  then  $r$  and  $g$  are in the same petal\* so the DDFS terminates. Otherwise  $a[g] < a[\alpha]$ , i.e.,  $g$  is in a petal\* of  $A$  preceding  $A_\epsilon$ . As in state Disjoint, move  $\alpha$  to catch up with  $g$ , and enter state Meet.

### 3.2.3 Finishing the DDFS

If an ap is discovered, delete it from the graph, and (as in [5]), repeatedly delete all vertices and petal\*'s with no predecessors. This deletes  $V(A)$ , by definition of anchored DDFS.

If a bottleneck  $b$  is discovered,  $b = \alpha = \text{bud}^*(A_\epsilon)$ . Merge petal\*'s  $A_\epsilon, \dots, A_\ell$ . In the data structure pop the  $\hat{B}$  stack to erase the buds above  $b$ , and add the new vertices of the petal\* to  $\hat{A}$ . Also merge appropriate bookkeeping lists into those for the new petal\*.

Now in either case start the next iteration.

### 3.3 Time Bound

An anchored DDFS is a special case of DDFS. So as in the Micali-Vazirani algorithm, all the anchored DDFSs use total time  $O(m)$  plus the time for set merging. It is easy to see we use  $O(1)$  time to identify the  $\text{bd}^*$  of a vertex in  $A$ . Thus all anchored DDFSs use total time  $O(m)$ .

Beyond DDFS the algorithm performs a number of dfs's to extend  $A$ . An edge of the graph gets scanned at most once in all these searches. Hence they use total time  $O(m)$ .

## References

- [1] H.N. Gabow, *Path-based depth-first search for strong and biconnected components*, Inf. Proc. Letters, 74 (2000), pp. 107–114.
- [2] H.N. Gabow and R.E. Tarjan, *A linear-time algorithm for a special case of disjoint set union*, J. Comp. and System Sci., 30, 2 (1985), pp. 209–221.
- [3] H.N. Gabow and R.E. Tarjan, *Faster scaling algorithms for general graph matching problems*, J. ACM, 38, 4 (1991), pp. 815–853.
- [4] S. Micali and V.V. Vazirani, *An  $O(\sqrt{|V|}|E|)$  algorithm for finding maximum matching in general graphs*, in Proc. 21st Annual IEEE Symp. on Foundations of Comp. Sci., 1980, pp. 17–27.
- [5] V.V. Vazirani, *An improved definition of blossoms and a simpler proof of the MV matching algorithm*, manuscript, April 24, 2013.